NIST Cloud Computing Reference Architecture

Contracts and SLA

DRAFT Recommendations of the National Institute of Standards and Technology (NIST) Public Working Group

Table of Contents

Cloud Contracts and Service Level Agreements (SLA)

Background

Contracts and service level agreements play a key role in the procurement of cloud computing services. An important caveat with cloud computing is that the consumer may have an agreement with one provider, but the service may be delivered via a myriad of subcontractors or other dependencies of the cloud provider and the consumer has no direct contractual relationship with these additional parties. The consumer, in fact, may have no knowledge of these third parties unless the provider chooses, or is otherwise required, to disclose them, and yet these entities may incur risk for which the consumer could ultimately be liable.

The default contract offered by cloud providers is often written to protect the provider's interests and can severely constrain the visibility of the consumer into the delivery mechanisms of the service - including dependencies on other providers and their suitability, reliability, and trustworthiness to handle consumer data. Additionally, the default contract typically limits options for the consumer in the case of failure to perform by the provider or a subcontractor thereof. It is therefore incumbent upon the cloud consumer to perform due diligence on potential cloud providers along with their dependency on partners or other providers to deliver the service, and be prepared to negotiate cloud service contracts which appropriately share the liability between consumer and provider. To facilitate negotiation requires a list of core concepts that should be discussed and agreed to between the parties to in order to create a high quality and meaningful SLA.

The Cloud Contracts/SLA team was tasked with identifying the common elements in existing cloud contracts and SLAs along with research published to date for the purpose of developing a taxonomy and identifying commonalities. It was noted that there were similarities in existing contracts and SLAs even though they did not frequently utilize the same structure or exact language. To define the key terms observed, the following steps were taken:

1. As this guidance is primarily for federal agencies, definitions in existing publications pertinent to the U.S. government were given preference
2. Where multiple definitions existed, especially for legal and IT terms, those most suited to the context of cloud computing contracts and SLAs were selected and/or modified as necessary

Contracts and SLAs examined were limited to those publicly available and thus do not necessarily reflect the structure or language of negotiated contracts between parties. While there may also be an automated component to SLAs, that portion should be derived from the negotiated contract and therefore is not specifically addressed in this work. Additionally, contracts and SLAs examined were for

parties in the role of consumer and provider. Brokerage functions and multi-party contracts were not examined in this work.

Scope

To fully understand the scope of contracts and service level agreements in cloud computing, it is necessary to examine the definitions of each:

> **Service agreement**: A legal document specifying the rules of the legal contract between the cloud user and the cloud provider. (NIST SP 800-146)
>
> **Service-level agreement**: A document stating the technical performance promises made by the cloud provider, how disputes are to be discovered and handled, and any remedies for performance failures. (NIST SP 800-146)

The service agreement - alternately known as master service agreement (MSA), terms of service (ToS), terms and conditions (T & C), or simply "the contract" –is the higher order document in agreements between parties and the service-level agreement (SLA) is subservient. This is an important distinction because the SLA acronym is frequently, and incorrectly, used to reference the contractual relationship as a whole – a role that an SLA alone is incapable of performing. The service agreement addresses the whole of the contractual relationship and typically contains – in addition to the SLA and terms of service – policies for acceptable use and privacy. Remedies for failure to perform are occasionally specified in the service agreement – arbitration is one such example - but parties may also have the right to litigation if breach of contract can be established.

The service-level agreement has, as its scope, the performance levels of the provider as agreed to by the consumer and typically the sole remedy for failure is the issuance of service credits by the provider. Whereas broader components of SLAs may be contained in the written portion of the document, service level agreements may also have an automated component to facilitate measurement and enforcement dynamically.

Need for Metrics

The term “metric” is not consistently defined within the Information and Communication Technology (ICT) Industry. For our purposed, the term “metric” is used to describe individual “measures” such as the number of users and “metrics” such as “Gigabytes per Second”.

Cloud metrics largely fall into two major categories: business metrics (often defined within the SLA) and technical metrics (monitoring metrics) that enable the business SLA to be met.

For example, “response time” may be specified in the SLA, meanwhile other technical measures such as “hops” and “bandwidth” may be used to dynamically allocate resources, enabling “response time” SLAs to be met. Usage based costing metrics are generally a sub-category of the business metrics and will be a major component of a Service Agreement or Service Level Agreement. Some examples of usage based metrics are: Number of Users, Instance Minutes, Storage Resource Capacity Used Bytes, CPU Minutes and RAM in Megabytes. Costs metrics are established based on dollars per unit (“$/ Instance Minute” for example).

SLA metrics require appropriate categorization, clarifying to align with SLA objectives and specify consequences when SLAs are not met.

Context	Response time for requesting / obtaining additional storage
Constraints	During business hours 00:00 GMT to 12:00 GMT (no guarantee outside of business hours)
Measures (1)	Date / Time of Request to Cloud Provider from Cloud Consumer (triggered by storage request)
Measures (2)	Date / Time of Successful Completion Response from Cloud Provider
Metrics Calculation	Measure (2) - Measure (1)
Collection Method	Automated, triggered as part of service request
Units	Milliseconds
Used & Consequences	Used: Cloud Provider guarantee is a maximum of 3000 Millisecond response time for IaaS Storage requests 00:00 to 12:00. For every 10 IaaS storage requests exceeding 3000 Millisecond response a 10 % reduction will be applied to total IaaS storage charges for the month

Table 1: Example “response time” metric analysis

Table 1 is a simplified example of how the metric of “response time” can be defined and assessed in context with SLAs, clarifying response time in a “real-life” scenario.

Metrics considerations are dependent on the supported service models (IaaS, PaaS and SaaS) and the type of services provided within that model, for example, network, storage and computing services for IaaS.

In the NIST Cloud SLA Taxonomy the "metrics" (response time, availability throughput) identified are examples of 3 frequently cited metrics in context of SLAs, however are not an inclusive list. Appropriate metrics may be identified around any node in the SLA taxonomy mind map. For example, "Recover" might use "Mean Time To Recover" as a metric as part of a SLA. The SLA mind map helps foster discussions, helping identify potential metrics that will help quantify, monitor and assure that prioritized SLA expectations are being met. Metrics in context of SLAs often have stated expectations related to minimums, maximums, defaults and consequences for deviations from stated objectives.

In summary, when considering metrics in a cloud SLA, it is recommended that consumers and providers:

- Understand the business objectives for the cloud opportunity.
- Understand context and where the stakeholders fit into the cloud ecosystem.
- Understand potential cascading SLAs and associated metrics.
- Understand enabling "technical metrics" vs more visible "business metrics".
- Identify the set of metrics that align with prioritized objectives.
- Understand the usage cost models that are applied
- Clarify how the metrics will be used and what decisions will be made
- Ensure these metrics are defined at the right level of granularity and can be monitored on a continuous basis.
- Determine available standards that help provide a consistent measurement method. (some will evolve as cloud computing matures)
- Understand the value and limitations of the metrics collected
- Analyze and leverage the metrics on an ongoing basis as a tool for influencing business decisions.

Cloud Computing metrics provides critical information to optimize cloud experiences, perform comparative analysis and help make informed decisions. Metrics are a cornerstone for transparent Service Level Agreement management and good governance.

Metrics are a key part of successful cloud service level agreements and will be covered in more detail in a future publication.

Requirements for Government Agencies

In contracting for cloud computing services, it is essential to consider that federal agencies have some unique obligations that may not be required in the private sector. Directives on records retention and maintenance, security, privacy, and public access place additional responsibilities on federal agencies that must be considered when procuring cloud services.

Compliance

In addition to laws and regulations that pertain to the private sector, government cloud consumers also have to comply with additional requirements such as:

- Computer Fraud and Abuse Act [PL 99-474, 18 USC 1030]
- E-Authentication Guidance for Federal Agencies [OMB M-04-04]
- Federal Information Security Management Act (FISMA) of 2002 [Title III, PL 107-347]
- Freedom of Information Act as Amended in 2002 [PL 104-232, 5 USC 552]
- Guidance on Inter-Agency Sharing of Personal Data – Protecting Personal Privacy [OMB M-01-05]
- Homeland Security Presidential Directive-7, Critical Infrastructure Identification, Prioritization, and Protection [HSPD-7]
- Internal Control Systems [OMB Circular A-123]
- Management of Federal Information Resources [OMB Circular A-130]
- Management's Responsibility for Internal Control [OMB Circular A-123, Revised 12/21/2004]
- Privacy Act of 1974 as amended [5 USC 552a]
- Protection of Sensitive Agency Information [OMB M-06-16]
- Records Management by Federal Agencies [44 USC 31]
- Rehabilitation Act of 1973 [Section 508 Amendment]
- Responsibilities for the Maintenance of Records About Individuals by Federal Agencies [OMB Circular A-108, as amended]
- Security of Federal Automated Information Systems [OMB Circular A-130, Appendix III]

The Federal Information Security Management Act of 2002 (FISMA) "requires agencies to provide information security protections commensurate with risks and their potential harms to government IT systems" (OMB, 2011, p. 6). Agency heads are required to submit annual reports that provide a comprehensive summary of FISMA compliance for their information and systems, and as of January 1, 2011, monthly data feeds on selected reporting areas is mandated.

Adherence to these regulations places additional obligations upon federal agencies to ensure that cloud service providers can deliver services in a compliant manner, the cloud service(s) can be utilized by the federal agency in a compliant manner, and that agreements between the parties clearly specify the requirements to be met and the obligations of each party in meeting them.

FedRAMP

"The Federal Risk and Authorization Management Program (FedRAMP) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. This approach uses a 'do once, use many times' framework that will save cost, time, and staff required to conduct redundant agency security assessments" (FedRAMP ConOps, 2012, p.2). FedRAMP allows agencies to used pre-approved authorization packages from CSPs, which are verified by an independent third party assessment organization (3PAO), in the procurement of cloud services.

As a part of the authorization package, CSPs are required to provide a Control Implementation Summary (CIS) which specifies the control origin (e.g. consumer, provider, shared, etc.) along with a control definition and the status of same (e.g. in place, partially implemented, planned, etc.). Federal agencies should scrutinize the CIS to ensure the controls specified are acceptable as delineated and that language in the CSP contract and/or SLA accurately reflects responsibilities as outlined in the CIS.

Cloud Service and Service Level Agreement Mind Maps

A cloud service-level agreement (CSLA) is part of a service contract in which a particular cloud service level is set between a Cloud Provider and a Cloud Consumer. For the Cloud Consumer, this contract deals with two separate but linked concepts: the contract agreement itself and its associated cloud computing service. The main contract, typically called Service Agreement or Master Terms of Service, deals with the general provisions of the contract while the CSLA identifies the particular cloud service to be provided to the consumer along with associated information.

While both the Master Service Agreement and the CSLA are linked, to assist the consumer the concepts associated with both have been split into two taxonomies and associated mind maps. Figure 1 displays the master the master terms of service concepts organized around the general provisions of the contract and contains no cloud computing specific concepts beyond a single element labeled Cloud Service Level Agreement.

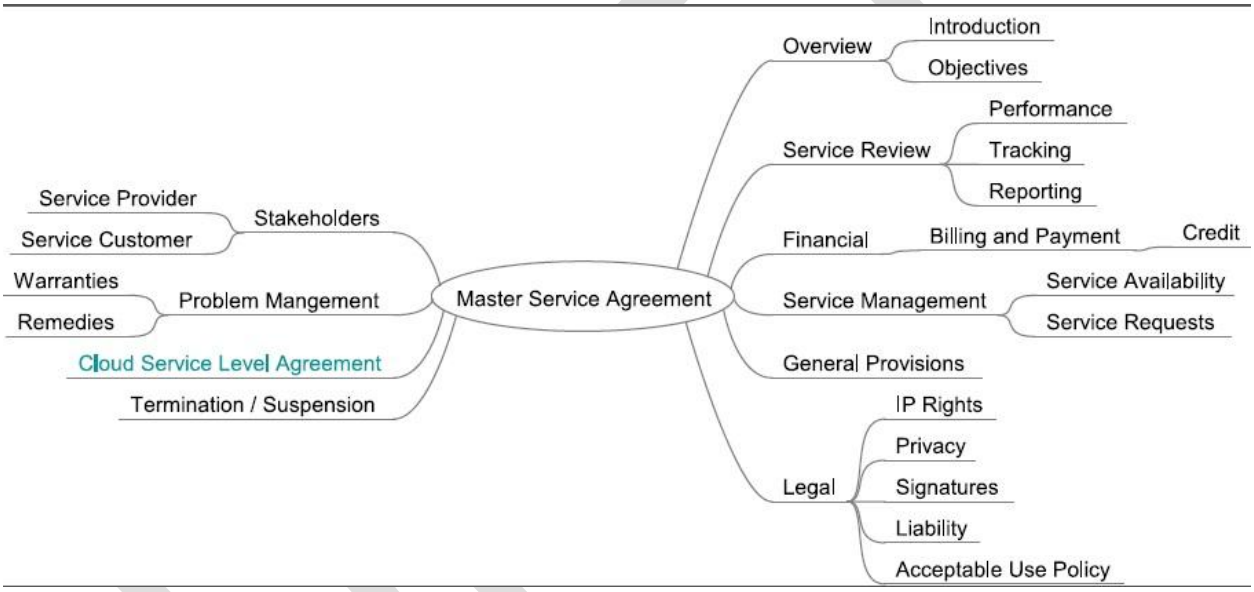

Figure 1. Mind Map of a Master Service Agreement for Cloud Computing

Compiled from various publicly available examples of master terms of service, research documents, and legal resources, the major elements of the Master Service Agreement are: Overview, Service Review, Financial, Service Management, General Provisions, Legal, stakeholders, problem management, and termination / suspension. While it is outside the scope of this document to delve deeply into each of these elements, the following information is provided on key areas of the mind map.

The Overview element typically contains an introduction into the purpose of the MSA and desired objectives. The Service Review element outlines how service delivery monitoring will occur within the contract and what parameters will be utilized. Three broad areas that should be specified are service

performance, tracking, and reporting. Performance for purposes of a cloud MSA can be defined as "a measure of what is achieved or delivered by a system, person, team, process, or IT service" (ITIL) and one instrument for determining performance is the SLA. However, the definition of what constitutes a failure to perform – and ultimately, a breach of the contract by the provider - may be defined in this element and should therefore receive careful scrutiny by consumers as some adverse situations are beyond the resolution capability of SLA remedies alone. Details of how tracking and reporting will occur and what obligations the provider has in these areas should be specified as well.

The Service Management element provides information about aspects of the service delivery, including availability and how requests are to be handled. The service requests portion can provide details on how such requests are initiated, by whom, what response times can be anticipated, and what obligations, if any, the provider has in responding to service requests. For service availability, it is worth examining a definition pertinent to the context of a cloud MSA:

> **Availability**: Specific service performance target - often expressed as a percentage (i.e. 99.99%) - that the cloud service provider will aim to meet (QMUL).

While the specifics of the availability as a performance target may be contained within the SLA , provider contracts may include a more broadly worded statement about the service being provided on an "as is" and "as available " basis, and such statements deserve careful inspection to ensure all parties understand what constitutes service availability versus what constitutes downtime – both planned and unplanned. It is also a consideration that, at some point, the issuance of service credits may not be a satisfactory remedy for significant and/or repeated failures of service availability on the part of the provider, whatever the cause.

The Legal element is where most of the legal clauses that appear in contracts are grouped. This grouping is for the convenience in analysis rather than an indication of where the actual clauses appear in contracts, as the formats vary considerably. The components listed here are not exhaustive, but representative of the most frequent items found in our study.

> **IP Rights** [Intellectual Property Rights over Service/Content]: Degree to which provider seeks to assert Intellectual Property rights over content and data uploaded to the cloud by customers. Additionally, a cloud provider may assert that the customer grants the provider a compulsory license to republish some or all of the customer's data for the purpose of provision of the service (QMUL).

While few of the cloud contracts examined attempted to assert any significant IP Rights over uploaded content for the types of services that government agencies would most likely utilize, the concept is presented here for informative purposes.

> **Privacy:** Information privacy is the assured, proper, and consistent collection, processing, communication, use and disposition of personal information (PI) and personally-identifiable information (PII) throughout its life cycle (OASIS).

Federal agencies have privacy responsibilities on many levels – Privacy Act of 1974, HIPAA Privacy Rule, OMB M-01-05, Privacy Impact Assessments, and privacy controls found in NIST 800-53 r4 to name a few. As such, appropriate contractual language is needed to ensure cloud service providers perform in a manner which facilitates compliance. As noted in SP 800-146 Draft, Legal Care of Subscriber Data also fits into the context of privacy by ensuring that providers do not expose data while in the process of providing the cloud service, or monitoring/logging the same. While data disclosure does not appear in the MSA mind map, it has privacy ramifications and so is included below.

> **Data Disclosure:** Terms of the circumstances in which providers will disclose customer information (including customer data stored on the provider's cloud). Providers will typically disclose such data in response to a valid court order, but some will accept requests (as distinct from enforceable orders) from recognized law-enforcement agencies, or where there is a clear and immediate need to disclose information in the public interest or to protect life. (QMUL)
>
> **Signatures:** A mark or sign made by an individual on an instrument or document, either hardcopy or electronic, to signify knowledge, approval, acceptance, or obligation (modified from FreeDictionary).

Liability is a complex legal concept such that three definitions were needed to provide enough clarity:

> **Liability (Direct):** Liability for damage caused to the customer by the provider. In this context, "direct liability‟ is taken to mean liability for losses to the customer relating to the loss or compromise of data hosted on the cloud service (QMUL).
>
> **Liability (Indirect):** A legal obligation resulting from damages awarded to an injured party because of the negligent act of a third party. In the context of cloud computing, typically used to cover indirect, consequential or economic losses arising from a breach by the provider (QMUL).
>
> **Liability Limits**: Terms seeking to limit the extent of any damages that the provider is held liable for.

Most off-the-shelf cloud service contracts examined in this study placed severe restrictions on liability of the cloud service provider. As offered, such services would be of limited use to federal agencies who could not realistically accept the potential risks involved.

> **Acceptable Use Policy:** A document that details the permitted (or in practice, forbidden) uses of the service (QMUL).

In the context of the Master Service Agreement, stakeholders are those parties who have an interest in the cloud service(s) being procured by the MSA. Typical stakeholders are cloud consumers and cloud providers, but there may be other parties -on which delivery of the service depends - who may not be clearly identified in the contract.

Problem Management is the process whereby the lifecycle of problems is managed, and within contracts two frequent vehicles are warranties and remedies. Below are the definitions of these terms, which often appear in contracts:

> **Warranty**: A promise or assurance that may be express, implied by the circumstances, or implied by law. A warranty might be an express or implied statement that particular facts are true (for example, that merchandise may be used for particular purposes). (Cornell)
>
> **Remedy**: The means by which a right is enforced or by which the violation of a right is prevented or compensated. (Lectric Law)

As mentioned previously, a frequent remedy for the violation of a service level agreement is the issuance of service credits by the provider. Most off-the-shelf cloud contracts explicitly disclaim the existence of any warranties, express or implied, and limit remedies to service credits (QMUL).

The legal section may also contain terms for whether the rights of the customer are affected due to the transfer of ownership of the provider (i.e., whether the customer or provider needs provide advance notice of changes or other provisions for changes for sale of the business), and expectations for insurance (property, safety, data breach, bonding, errors and omissions, etc.).

Figure 2 is an expanded view of the Cloud Service Level Agreement that appears in Figure 1 and was generated after researching publicly available CSLAs, research papers, position papers, policy statements, and acquisition guides. Care was taken to identify major elements that consistently appeared in multiple sources which indicated a high degree of relevance to the effort.

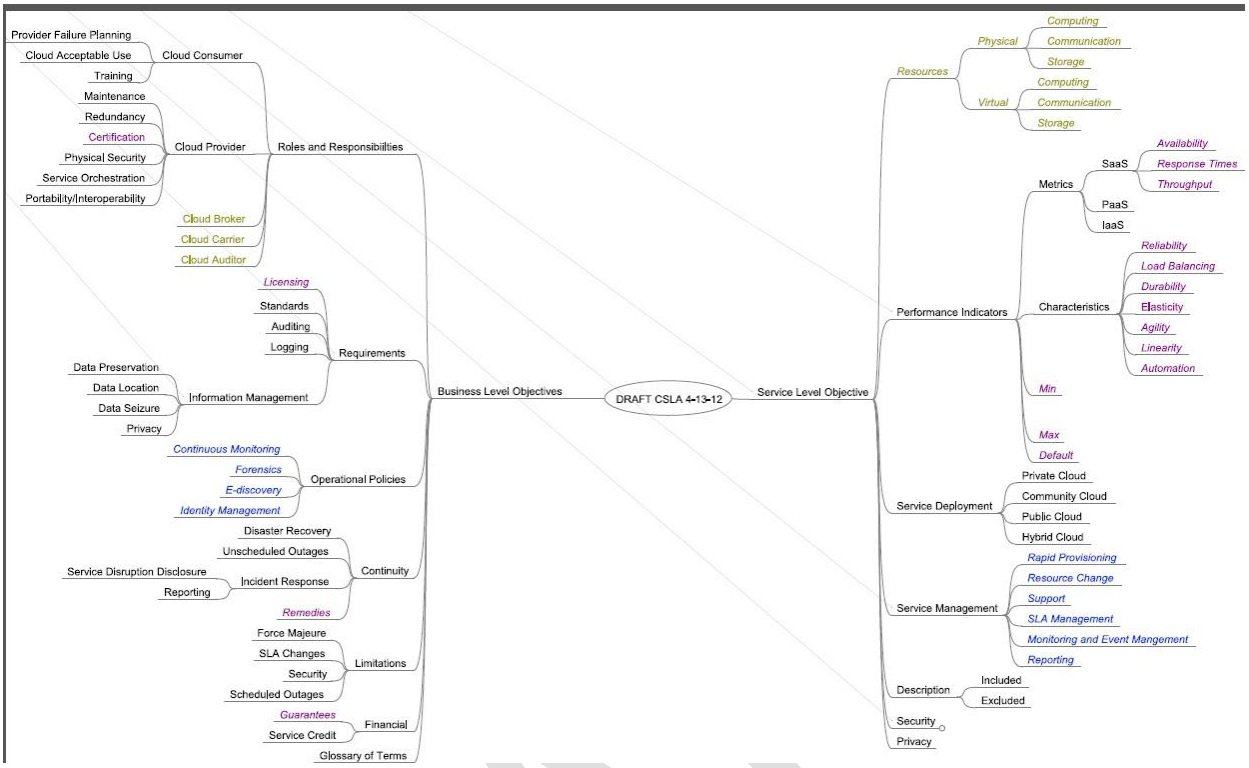

Figure 2. Mind Map of Service Level Agreement for Cloud Computing

Once the concepts were identified it was clear that at the topmost level there was a separation of concepts dealing with the business level objectives and service level objectives each with their own set of major sub concepts as shown in Figure 3.

Figure 3. Business vs Service Level Objectives

Within the Business Level Objectives:

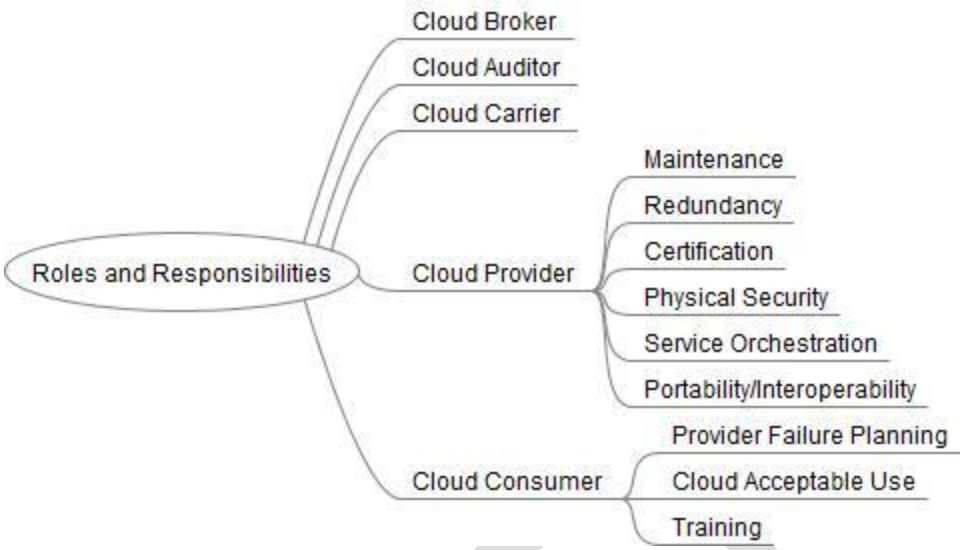

Figure 4. Major Stakeholder Roles and their Responsibilities in the Service Contract

Within the Cloud Service Agreement should be an identification of the major stakeholders within the contract and a detailed list of their responsibilities.

> **Roles and Responsibilities:** The delineation of titles and corresponding duties in the risk management framework as presented in NIST 800-37.

The major stakeholders are identified as the cloud computing roles identified within the NIST SP500 292 Cloud Computing Reference Architecture. Typical responsibilities associated with Cloud Provider include maintenance activities, planning for redundancy, any sort of certification activities, providing physical security for cloud resources, engaging in service orchestration, and providing support for portability and interoperability.

> **Certification:** The comprehensive evaluation of the technical and nontechnical security features of an AIS [automated information system] and other safeguards, made in support of the accreditation process, to establish the extent to which a particular design and implementation meets a set of specified security requirements. (FS 1037C)

> **Service Orchestration:** Refers to the arrangement, coordination and management of cloud infrastructure to provide different cloud services to meet IT and business requirements.

> **Portability:** The ability to transfer data from one system to another without being required to recreate or reenter data descriptions or to modify significantly the application being transported. (FS 1037C)

> **Interoperability:** The capability to communicate, execute programs, or transfer data among various functional units under specified conditions. (ANSDIT)

Typical Cloud Consumer responsibilities include performing provider failure planning, adhering to established acceptable use policy, and training the cloud users.

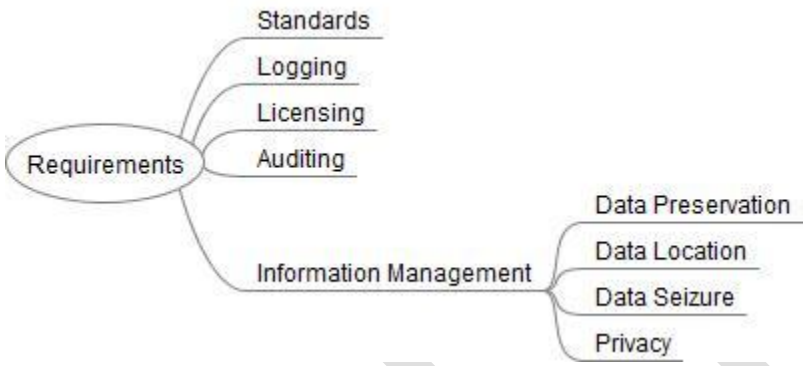

Figure 5.Cloud Business Requirements identified in the Service Contract

Within the Cloud Service Agreement will be a set of requirements that will be established by the cloud consumer and that must be followed by the cloud provider. Examples include adherence to a given set of standards, logging requirements, licensing requirements, information on how audits will be accommodated (possibly a right to audit clause), and how cloud consumer information will be managed. Information management can be quite detailed with requirements concerning how data will be preserved, where data can be physically located, what will happened if data is seized and how to handle privacy.

> **Data Preservation:** Disposition of customer data after the relationship with a cloud provider ends. Two issues: whether there will be any opportunity for the customer to gain access to the data (for example, to retrieve it for use elsewhere) once the contract has ended, and whether there is assurance from the provider that data will effectively be deleted after this stage. (QMUL)
>
> **Data Location:** The possibility that a customer's data may be stored or processed in a totally different, and potentially unknown, jurisdiction. (QMUL)
>
> **Data Seizure:** Seizure of property (data, in this instance) occurs when government action meaningfully interferes with an individual's (data owner, in this instance) possessory interest in that property [SCOTUS as referenced in Kerr].

Typically, the term data seizure references a government action involving a subpoena or warrant issued to a provider. Subsequently the target assets - the data, and possibly also the container(s) housing it -

are seized. The provider may have no obligation to notify the data owner, and in some instances may legally be prohibited from doing so.

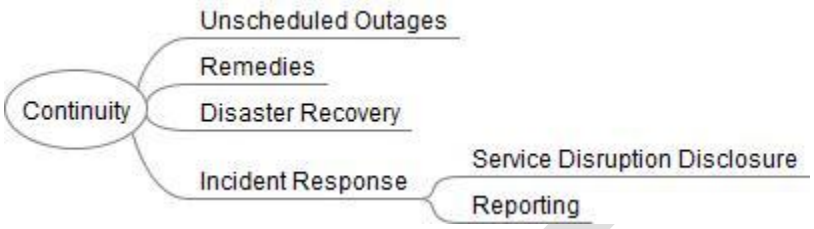

Figure 6.Business Continuity Issues

Unscheduled Outage [Disruption]: An unplanned event that causes the general system or major application to be inoperable for an unacceptable length of time (e.g., minor or extended power outage, extended unavailable network, or equipment or facility damage or destruction). (CNSSI 4009)

Remedies for unscheduled outages are typically limited to the issuance of service credits by the provider unless other recourse has been negotiated.

Disaster Recovery: The ability of an enterprise to maintain operational service levels in response to a major loss of enterprise capability or damage to its facilities. (Adapted from CNSSI 4009). There should be pertinent details about any joint testing needed to verify DR functionality.

Incident Response: A predetermined set of instructions or procedures to detect, respond to, and limit consequences of an incident against an organization's IT systems(s). (Modified from CNSSI 4009)

Service Disruption Disclosure: Typically defines the conditions (severity, timeframe, etc.) under which a provider must disclose the details of a service disruption. The term may also define what constitutes a service disruption and what, if any, obligation to notify consumers exists.

Reporting: Determines how the CSP will report incidents, and conversely, how the consumer will report incidents to the provider.

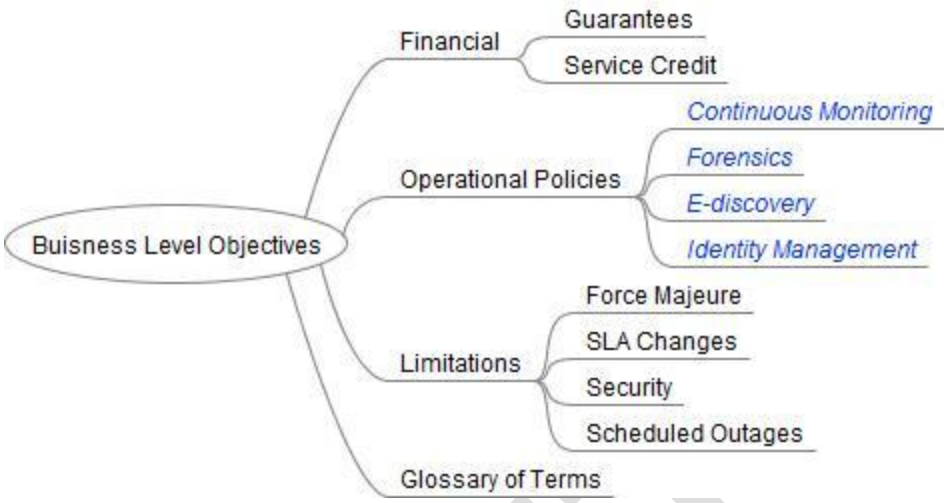

Figure 7.Financial, Operational Policies, Limitations, and General Glossary of Terms Elements

Financial

Guarantees: In the context of cloud services, this typically refers to a promise to make a product good if it has some defect. It may also/instead refer to the promise to achieve or maintain a specific performance objective. (Modified from Law.com)

Service Credit: A rebate against future billing awarded to customers as compensation for failure to deliver the service to set levels. (QMUL)

Operational Policies

Continuous Monitoring: The process implemented to maintain a current security status for one or more information systems or for the entire suite of information systems on which the operational mission of the enterprise depends. (CNSSI 4009)

Forensics: The scientific examination and analysis of data held on, or retrieved from, ESI in such a way that the information can be used as evidence in a court of law. It may include the secure collection of computer data; the examination of suspect data to determine details such as origin and content; the presentation of computer based information to courts of law; and the application of a country's laws to computer practice. Forensics may involve recreating "deleted" or missing files from hard drives, validating dates and logged in authors/editors of documents, and certifying key elements of documents and/or hardware for legal purposes. (Sedona)

Electronic Discovery: The process of identifying, preserving, collecting, preparing, reviewing, and producing electronically stored information ("ESI") in the context of the legal process. (Sedona)

Identity Management: a broad administrative area that deals with identifying individuals in a system (such as a country, a network, or an enterprise) and controlling their access to resources

within that system by associating user rights and restrictions with the established identity. (TechTarget)

The Limitations element covers areas where cloud providers denote the limitations of the service offering and their liability.

Force Majeure: Clauses which excuse a party from liability if some unforeseen event beyond the control of that party prevents it from performing its obligations under the contract. (Yale Law Library)

SLA Changes: Terms which specify how and when the provider may change the terms of the SLA and whether or not there is any obligation to notify the customer(s) affected. (QMUL)

Security: It is incumbent upon cloud consumers to read carefully the de facto security included as a part of the service offering. Providers frequently disclaim any significant responsibility for security of user information created, transported, or processed within their services.

Scheduled Outages: Also referenced as Planned Downtime, Scheduled Maintenance, et al. These are service disruptions initiated by the provider to undertake system maintenance and/or upgrades. This type of outage is typically excluded from remedies offered to specified unscheduled or unplanned disruptions.

A glossary of terms should be included within the contract to provide clear explanations for key terms which appear in the document, as multiple definitions or perceptions may exist.

Within the Service Level Objectives:

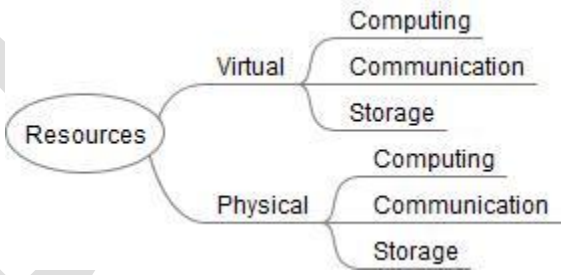

Figure 8.Resources Identified within the Contract (both Virtual and Physical)

Resources: With Cloud Computing the providers computing resources are generally pooled to service multiple consumers using a multi-tenant model. With different physical and virtual resources dynamically assigned and reassigned according to consumer demand. The SLA should specify specific resources and types of resources required by the contract.

Virtual: creating a virtual version of a device or resource such as storage, processing, memory, and network bandwidth**.**

Physical: the actual device or resource such as storage, processing, memory, and network bandwidth**.**

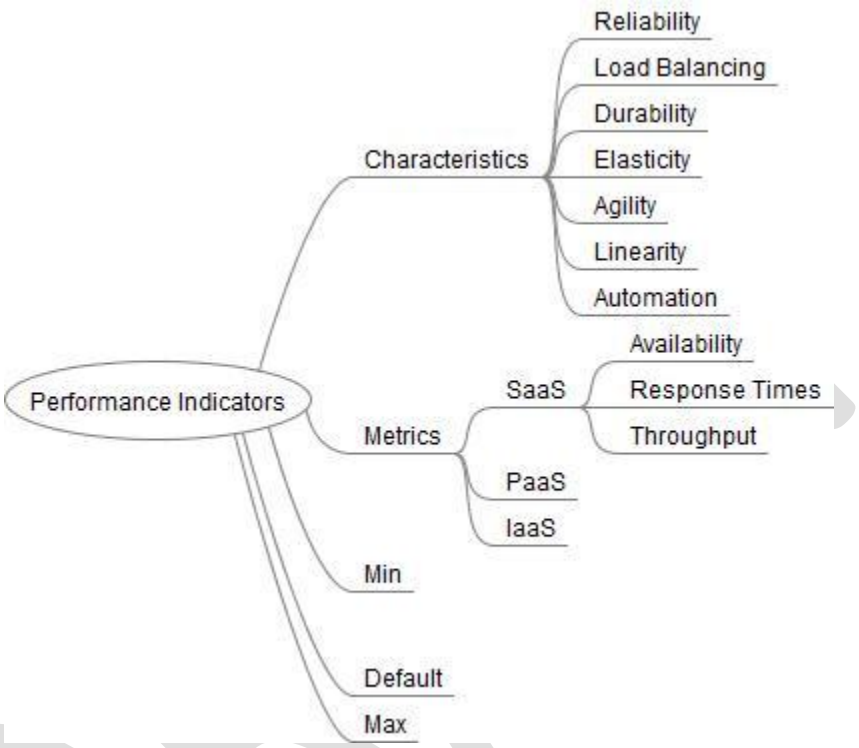

Figure 9.Performance Indicators that will be used to measure the service agreement

Performance Indicators:

Characteristics: qualities that define or describe a cloud service. Common characteristics in use within SLAS include: reliability, load balancing, durability, elasticity, agility, linearity, automation.

Metrics: A method of measuring something related to a cloud computing service. Quite often the metrics used for any specific SLA are dependent on the type of cloud computing service that is being used by the consumer. For example, for Software as a Service, metrics that could be selected include availability, response time, or throughput.

Min: The minimum amount or setting for a cloud service performance indicator

Default: A preselected option adopted for the performance indicator when no alternative is specified by the cloud consumer

Max: A maximum amount or setting for a cloud service performance indicator

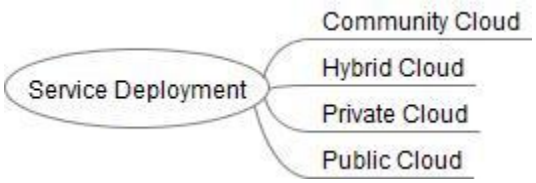

Figure 10.How the Service will be Deployed for the Consumer

Service Deployment refers to all the activities and organization needed to make a cloud service available to a consumer. (NIST SP500-292). Within the context of a SLA, the cloud provider will create a service for the consumer in one of four possible deployment configurations including:

Community Cloud: The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises. (NIST SP800 – 145)

Hybrid Cloud: The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds). (NIST SP800 – 145)

Private Cloud: The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises. (NIST SP800 – 145)

Public Cloud: The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider. (NIST SP800 – 145)

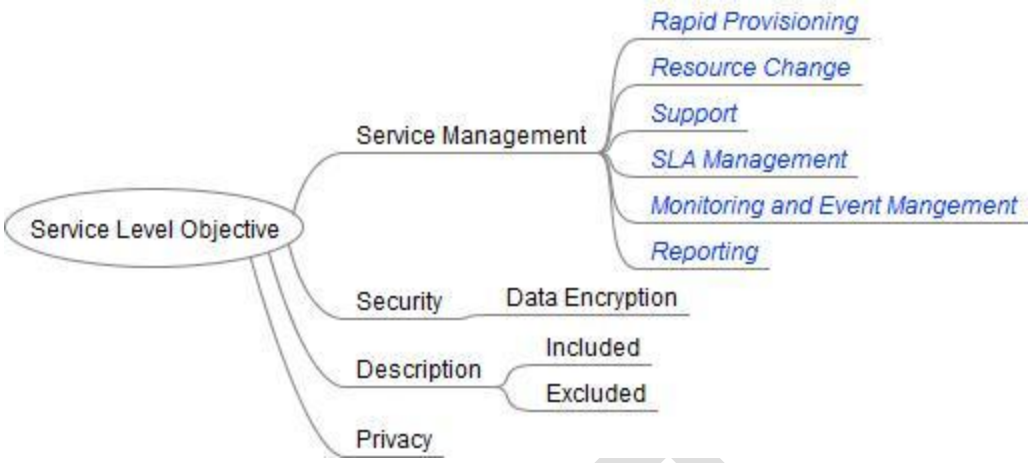

Figure 11.Service Management, Security, Privacy, and Description of Services

Service Management refers to includes all the service-related functions that are necessary for the management and operations of those services required by or proposed to customers. (NIST SP 500 – 292). Common functionality addressed within service management include: rapid provisioning, resource change, support, SLA management, monitoring / event management, and report.

Security: Identifies the security mechanisms and processes that are to be included within the service identified within the contract. While this is typically up for negotiation between the consumer and provider, the mandated use of data encryption is common. The consumer should also reconcile security specifications in the SLA with any limitations noted in the MSA.

Description: Typically describes the services that are specified within the contract with a focus on what elements are included within the specified service and what is excluded.

Privacy: Identifies how information privacy is contractually established within the SLA where Information privacy is the assured, proper, and consistent collection, processing, communication, use and disposition of disposition of personal information (PI) and personally-identifiable information (PII) throughout its life cycle. (Source: adapted from OASIS)

Conclusions

Contracts and service level agreements are key components in cloud computing services, but these are arguably the least understood and there are no broadly accepted standard frameworks or language for them (except as noted in the automated agreement portions where they exist) which fit within the required scope and context for cloud service procurement by federal agencies. While it is presumed that initial cloud service contracts will primarily involve the consumer and provider roles, multiple instances of these roles may certainly occur along with the possible addition of one or more brokerage roles. However, the basic tenets of establishing fair cloud computing contracts and SLAs are unchanged: the apportionment of responsibilities and accompanying risks among the parties involved along with clearly defined specifications and methods for ascertaining performance.

As cloud computing services evolve, the vehicles which form the basis of the relationship between the roles (i.e. contracts and SLAs) will need to evolve as well. The terms and their definitions as contained in this document may need to be modified, or other definitions added, in order to better reflect the context and usage in cloud computing service agreements. However, they form a basis for facilitating meaningful discussion between cloud service providers and cloud service consumers in creating a high quality and useful SLA.